09/743268

PATENT COOPERATION TREATY

From the INTERNATIONAL BUREAU

To:

Commissioner
US Department of Commerce
United States Patent and Trademark Office, PCT
2011 South Clark Place Room CP2/5C24
Arlington, VA 22202
ETATS-UNIS D'AMERIQUE

in its capacity as elected Office

RECEIVED
MAY 4 - 2001
Technology Center 2100

PCT

NOTIFICATION OF ELECTION

(PCT Rule 61.2)

Date of mailing (day/month/year) 06 April 2001 (06.04.01)	
International application No. PCT/NO00/00155	**Applicant's or agent's file reference** Opti44Fast7
International filing date (day/month/year) 10 May 2000 (10.05.00)	**Priority date** (day/month/year) 10 May 1999 (10.05.99)
Applicant SVINGEN, Børge et al	

1. The designated Office is hereby notified of its election made:

   [X] in the demand filed with the International Preliminary Examining Authority on:

   11 December 2000 (11.12.00)

   [ ] in a notice effecting later election filed with the International Bureau on:

2. The election [X] was
   [ ] was not

   made before the expiration of 19 months from the priority date or, where Rule 32 applies, within the time limit under Rule 32.2(b).

The International Bureau of WIPO **34, chemin des Colombettes** **1211 Geneva 20, Switzerland** Facsimile No.: (41-22) 740.14.35	Authorized officer Nestor Santesso Telephone No.: (41-22) 338.83.38

K, 09/723268 (2FC1 [illegible]).

14

REC'D 27 AUG 2001
WIPO PCT

PATENT COOPERATION TREATY

PCT

INTERNATIONAL PRELIMINARY EXAMINATION REPORT

(PCT Article 36 and Rule 70)

Applicant's or agent's file reference Opti44Fast7	**FOR FURTHER ACTION** See Notification of Transmittal of International Preliminary Examination Report (Form PCT/IPEA/416)	
International application No. PCT/NO00/00155	International filing date *(day/month/year)* 10/05/2000	Priority date *(day/month/year)* 10/05/1999
International Patent Classification (IPC) or national classification and IPC G06F17/30		
Applicant FAST SEARCH & TRANSFER ASA ET AL.		

1. This international preliminary examination report has been prepared by this International Preliminary Examining Authority and is transmitted to the applicant according to Article 36.

2. This REPORT consists of a total of 7 sheets, including this cover sheet.

   ☒ This report is also accompanied by ANNEXES, i.e. sheets of the description, claims and/or drawings which have been amended and are the basis for this report and/or sheets containing rectifications made before this Authority (see Rule 70.16 and Section 607 of the Administrative Instructions under the PCT).

   These annexes consist of a total of 14 sheets.

3. This report contains indications relating to the following items:

   - I ☒ Basis of the report
   - II ☐ Priority
   - III ☐ Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
   - IV ☐ Lack of unity of invention
   - V ☒ Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations suporting such statement
   - VI ☐ Certain documents cited
   - VII ☐ Certain defects in the international application
   - VIII ☐ Certain observations on the international application

Date of submission of the demand 11/12/2000	Date of completion of this report 24.08.2001
Name and mailing address of the international preliminary examining authority: European Patent Office D-80298 Munich Tel. +49 89 2399 - 0 Tx: 523656 epmu d Fax: +49 89 2399 - 4465	Authorized officer Schneider, M Telephone No. +49 89 2399 7509

INTERNATIONAL PRELIMINARY EXAMINATION REPORT

International application No. PCT/NO00/00155

I. Basis of the report

1. With regard to the **elements** of the international application *(Replacement sheets which have been furnished to the receiving Office in response to an invitation under Article 14 are referred to in this report as "originally filed" and are not annexed to this report since they do not contain amendments (Rules 70.16 and 70.17))*:
   Description, pages:

1-11	as received on	16/07/2001	with letter of	11/07/2001

   Claims, No.:

1-9	as received on	16/07/2001	with letter of	11/07/2001

   Drawings, sheets:

1-6	as originally filed

2. With regard to the **language**, all the elements marked above were available or furnished to this Authority in the language in which the international application was filed, unless otherwise indicated under this item.

   These elements were available or furnished to this Authority in the following language: , which is:

   ☐ the language of a translation furnished for the purposes of the international search (under Rule 23.1(b)).

   ☐ the language of publication of the international application (under Rule 48.3(b)).

   ☐ the language of a translation furnished for the purposes of international preliminary examination (under Rule 55.2 and/or 55.3).

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, the international preliminary examination was carried out on the basis of the sequence listing:

   ☐ contained in the international application in written form.

   ☐ filed together with the international application in computer readable form.

   ☐ furnished subsequently to this Authority in written form.

   ☐ furnished subsequently to this Authority in computer readable form.

   ☐ The statement that the subsequently furnished written sequence listing does not go beyond the disclosure in the international application as filed has been furnished.

   ☐ The statement that the information recorded in computer readable form is identical to the written sequence listing has been furnished.

4. The amendments have resulted in the cancellation of:

   ☐ the description, pages:

   ☐ the claims, Nos.:

☐ the drawings, sheets:

5. ☒ This report has been established as if (some of) the amendments had not been made, since they have been considered to go beyond the disclosure as filed (Rule 70.2(c)):

   (Any replacement sheet containing such amendments must be referred to under item 1 and annexed to this report.)
   see separate sheet

6. Additional observations, if necessary:

V. Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	1-9
	No:	Claims	
Inventive step (IS)	Yes:	Claims	1-9
	No:	Claims	
Industrial applicability (IA)	Yes:	Claims	1-9
	No:	Claims	

2. Citations and explanations
   see separate sheet

Re Item I

Basis of the report

(i) The feature of the amended claim 1 that

the number of search nodes (N_B) in each search node column (S) scales *with estimated or expected traffic load*, thus implementing a scaling of performance,

although being obvious, is not disclosed in the application as originally filed, where only is said (s. originally filed claim 1, last paragraph, and claim 10) that

a scaling of performance takes place through replication of one or more partitions d, being implemented by an increase of the number of search nodes (N_B) in each search node column (S).

(ii) The feature that

scaling of performance takes place through replication of one or more partitions d,

as is explicitly stated in the originally filed claim 1 and is valid for the whole application as originally filed, is no more explicitly stated in the amended claim 1, thus allowing the interpretation that

in any case searching the document collection D takes place with a single search node (N_B) *holding by reference* one of the partition-dependent data sets $d_{p,k}$ and all search nodes (N_B) in a search node column (S_B) *containing by reference* identical partition-dependent data sets $d_{p,k}$,

which goes beyond the disclosure of the application as originally filed.

Re Item V

Reasoned statement under Art. 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

Reference is made to the following document:

D1: Erhard Rahm, Mehrrechner-Datenbanksysteme, Addison-Wesley, 1994, p. 34,35,62-66,326-328

The document D1was not cited in the international search report. A copy of the document was sent to the applicant.

1. The subject-matter of claim 1 is new and inventive in the sense of Articles 33(2) and 33(3) PCT.

State of the Art

It is straightforward to implement an Internet search machine by setting up a database relation, in which each tupel comprises a keyword and the URL of a document containing the keyword. For storage and performance reasons it may be necessary to fragment the relation and to assign the fragments to different workstations, each corresponding to a search node column of the application, in a parallel shared-nothing database system (see D1, page 34, lines 17-24, page 35, lines 7-24). In order to have concurrent execution of requests as "return all URLs of documents containing keyword 1 and keyword 2" with minimum communication among the workstations, it is obvious to

horizontally fragment in respect to ranges or hash function results of the URLs, i.e. each fragment, and thus workstation, corresponds to a particular partition of the URLs, or documents (see D1, page 326-page 328, line 22, chapters 5.3.1 and 5.3.3).

In this system

> the number v of search node columns (S) scales with the number n of partitions d, thus implementing a scaling of the data volume.

It may happen that a maximum number of partition or fragment ranges, that is feasible, is reached or that a fragmentation may not be changed easily. In such a situation, in order to further improve performance it is the most obvious to realize the workstations to which the partitions are assigned as high-end multiprocessor workstations which already are used for Internet search engines, whereby each processor corresponds to a search node of the application. Then,

> a scaling of performance is implemented by an increase of the number of search nodes (N_B), i.e. processors, in each search node column (S), i.e. multiprocessor workstation,

but this takes place without replication of one or more partitions d. Thus, in this case searching the document collection D takes place with a single search node (N_B) *merely having access to, i.e holding by reference only,* the partition-dependent data set $d_{p,k}$ stored on the respective multiprocessor workstation, i.e. all search nodes, or processors, in a search node column (S_B) *are containing by reference only* identical partition-dependent data sets $d_{p,k}$.

Problem

The use of high-end multiprocessor workstations may be very expensive, and they are scaleable only within limits.

Solution

The application proposes that

> in any case searching the document collection D takes place with a single search node (N_B) holding one of the partition-dependent data sets $d_{p,k}$ and all search nodes (N_B) in a search node column (S_B) containing identical partition-dependent data sets $d_{p,k}$,

whereby partitions are replicated, as is explicitly stated in the originally filed claim 1.

This allows alternative architectures as e.g. to realize search nodes by modules that hold replicated partitions on local memory, and hereby exploits the fact that no communication between search nodes is necessary.

2. Claims 2-9 are dependent on claim 1 and as such also meet the requirements of the PCT with respect to novelty and inventive step.

A search engine with two-dimensional linearly scalable parallel architecture.

The present invention concerns a search engine with two-dimensional linearly scalable parallel architecture for searching a collection of text documents D, wherein the documents can be divided into a number of partitions $d_1, d_2, \ldots d_n$, wherein the collection of documents D is preprocessed in a text filtration system such that a preprocessed document collection D_p is obtained and corresponding preprocessed partitions $d_{p1}, d_{p2}, \ldots d_{pn}$, wherein an index I can be generated from the document collection D such that for each previous preprocessed partition $d_{p1}, d_{p2}, \ldots d_{pn}$ a corresponding index $i_1, i_2, \ldots i_n$ is obtained, wherein searching a partition d of the document collection D takes place with a partition-dependent data set $d_{p,k}$ comprising both the preprocessed partition d_{pk} and the corresponding index i_k, with $1 \leq k \leq n$, and wherein the search engine comprises data processing units which form sets of nodes connected in a network.

Most prior art search engines work with large data set and employ powerful computers to perform the search. However, searching is a partitionable data processing problem, and this fact can be used to partition a search problem into a large number of specific queries and let each query be processed simultaneously on a commensurate number of processors connected in parallel in a network. Particularly searching can be regarded as a binary partitionable data processing problem, and hence a binary tree network is used for establishing a multiprocessor architecture such as disclosed for instance in US patent No. 4860201 (Stolfo & al.) and international patent application PCT/NO99/00308 which belongs to the applicant and hereby is incorporated by reference. The present applicant has developed proprietary technologies for searching within regular text documents. These technologies are i.a. based on a search system and a method for searching as described in international patent application PCT/NO99/00233 which belongs to the applicant and hereby is incorporated by reference. The search system is based on efficient core search algorithms which may be used in the search engine according to the invention.

However, it has become increasingly important to cater for a growing number of documents to be searched and also to be able to handle an increased traffic load, i.e. the number of queries per second which shall be processed by the

search system. This, apart from the ability to handle a large number of queries simultaneously on processor level, implies that a search engine should be implemented with an architecture that allows for preferably linear scalability in two dimensions, viz. both with regard to the data volume and the performance, i.e. the ability to handle a very large number of queries per second. Considering the development of the World Wide Web, a scalability problem in the search engine architecture will be extremely important as there presently is an enormous growth rate in both the number of documents and the number of users on the Internet.

Prior art search engine solutions for Internet are able to scale to a certain level, but almost always this is achieved in a manner that requires a high cost increase of the search engine system relative to the growth in data volume or data traffic. Very often the system costs scale as the square of the data volume or the traffic, a doubling of the data volume thus leading to quadrupled system costs. Furthermore all the major Internet search engines presently are based on very expensive server technology, often coupled with brute computing force-approaches and accompanied with disadvantages such as slow server turnaround, requirements for special hardware to provide fault tolerance etc. The system costs can e.g. be measured as the amount of hardware required to implement a search engine solution or the actual aggregated price of the system.

It is thus a main object of the present invention to provide a search engine with a multilevel data and functional parallelism, such that large volumes of data can be searched efficiently and very fast by a large number of users simultaneously.

Particularly it is a further object of the invention to provide a parallel architecture for implementing a search engine with a multilevel data and functional parallelism.

Yet a further object of the present invention is to provide a parallel architecture which is linearly scalable in two dimensions, i.e. with regard to both data volume and performance, that is the query rate.

The above-mentioned objects and further features and advantages are provided with a search engine according to the invention which is characterized in that the first set of nodes comprises *a* dispatch nodes, a

second set of nodes comprises b search nodes, a third set of nodes comprises g indexing nodes, and an optional fourth set of nodes comprises e acquisition nodes, that the dispatch nodes are connected in a multilevel configuration in the network, that the search nodes are grouped in v columns which are connected in parallel in the network between the dispatch nodes and an indexing node, that the dispatch nodes are adapted to process search queries and search answers, the search queries being dispatched further to all search nodes and in case the acquisition nodes are not present, the search answers being returned to the dispatch nodes and therein being combined to a final search result, that the search nodes each are adapted to contain search software, that the indexing nodes are adapted for generally generating indexes i for the search software and optionally for generating partition-dependent data sets $d_{p,k}$ to search nodes which comprise a search processor module, that in case acquisition nodes are present, these are connected in a multilevel configuration in the network similar to that of the dispatch node, and adapted for gathering answers to search queries and outputting a final result thereof, thus relieving the dispatch nodes of this task, that the number v of search node columns scales with the number n of partitions d, thus implementing a scaling of the data volume, and that the number of search nodes in each search node column scales with an estimated or expected traffic load, thus implementing a scaling of performance, whereby in any case searching the document collection D takes place with a single search node holding one of the partition-dependent data sets $d_{p,k}$ and all search nodes in a search node column containing identical partition-dependent data sets $d_{p,k}$.

According to the invention are advantageously the multilevel configuration of the dispatch nodes and the optional acquisition nodes network provided by hierarchical tree structures, and the multilevel configuration of the optional acquisition nodes is then preferably a mirror image of the multilevel configuration of the dispatch nodes, the hierarchical tree structures preferably being binary tree structures.

According to the invention each of the search nodes advantageously comprises a search software module.

Further, according to the invention at least some of the search nodes comprises at least one dedicated search processor module, each dedicated

search processor module being realized with one or more dedicated processor chips, which is adapted for parallel handling of a number of search queries. In this connection it is preferred that the dedicated search processor chips are provided in the search processor modules in y processor groups, each with z search processor chips and being connected with and adapted to receive data from a memory assigned to the processor group.

According to the invention the increase in the number of partitions can preferably be accompanied by a corresponding increase in the number of dispatch nodes and, in case, also in the number of acquisition nodes, and optionally also by an increase in the number of index nodes.

Finally it is according to the invention advantageous that the separate node sets each is implemented over one or more workstations connected in a data communications network.

The search engine according to the invention shall now be described in terms of non-limiting exemplary embodiments and with reference to the accompanying drawings, in which

fig. 1 shows a general overview of the architecture in a first embodiment of the search engine according to the invention,

fig. 2 a general overview of the architecture in a second embodiment of the search engine according to the invention,

fig. 3 schematically a search node with search software and dedicated hardware search chips,

fig. 4 the layout of a single module of search chips,

fig. 5 schematically the query handling in a dedicated search chip,

fig. 6 schematically the principle of performance scaling,

fig. 7 a multilevel dispatch node configuration based on a binary tree structure,

fig. 8 an overview of the architecture of a first embodiment of the search engine according to the invention, with an indication of the principle for two-dimensional scaling, and

fig. 9 an overview of the architecture of a second embodiment of the search engine according to the invention, with an indication of the principle for two-dimensional scaling.

Searching a large collection of independent documents is a highly parallel task. The search engine according to the invention employs parallelism on different levels as shall be discussed in the following.

The search engine according to the present invention searches a document collection of documents D. The documents can be divided into n partitions $d_1, d_2, \ldots, d_n$. Each document collection D, or partition d of a document collection can be preprocessed for use in a hardware text filtering system, for instance implemented by dedicated hardware like the applicant's so-called Pattern Matching Chip (PMC) which is disclosed in the applicant's international patent application No. PCT/NO99/00344 which hereby is incorporated by reference. The preprocessed document collection is denoted D_p and the corresponding preprocessed document collection partitions d_p are denoted $d_{p1}, d_{p2}, \ldots, d_{pn}$.

Software-based search systems require an index generated from the document collection. The index is denoted I and the indexes corresponding to the document collection partitions $d_{p1}, d_{p2}, \ldots d_{pn}$, are denoted $i_1, i_2, \ldots i_n$.

The data set needed for searching a partition d of the document collection D, is called the partition-required or partition-dependent data set. In a software only system (SW system) the data set is the index i_k, while in the systems with hardware (SW/HW systems), the data set also includes the preprocessed document collection partition $d_{p,k}$ with the corresponding index i_k, where $1 \leq k \leq n$.

The essentially software-implemented partitioning and preprocessing operations can be rendered schematically as $*(D) \rightarrow *(d_1, \ldots, d_n) \rightarrow *(d_{p1}, \ldots, d_{pn}) \rightarrow d_{pk}$, where $*(D)$ denotes a partitioning operation on the input D, $*(d_1, \ldots, d_n)$ a filtering operation. e.g. indexing, on $d_1, \ldots, d_n$ and $d_{p,k}$ of course is the partition-dependent data set, which in an SW system only shall be the index i_k, and with $1 \leq k \leq n$.

A search engine is implemented on a cluster of workstations that are connected using a high performance interconnect bus. The not shown workstations then constitute the server of the search system (search server).

The workstations implement nodes of the search server. The nodes perform different tasks and are according to the invention implemented as set out below.

The nodes can be regarded as virtual nodes distributed among the workstations, but in a SW/HW search engine the dedicated search processing hardware must be physically present in some workstations in order to support the hardware-based search nodes. The search node software may then still be distributed. Also, some search nodes in SW/HW search engine may comprise software only, and optionally be distributed over more than one workstation.

- Dispatch nodes $N_{\alpha 1}$,...$N_{\alpha a}$ handle incoming queries and send queries along to all the search nodes. The dispatch nodes can also be configured as acquisition nodes, gathering the answers, i.e. the search results to the queries. Upon receiving the answers, the dispatch nodes in the acquisition mode merge the search results into a final result.
- Search nodes $N_{\beta 1}$,...,$N_{\beta b}$ hold a portion of the entire data set d_p. A search node comprises both the dedicated search software as well as a number of the above-mentioned PMC modules for hardware searching.
- Indexing nodes $N_{\gamma 1}$,...,$N_{\gamma g}$ are responsible for generating indexes for the dedicated search software on a number of search nodes. For the PMC modules, the indexing node N_γ also filters and preprocesses the raw data.
- Acquisition nodes $N_{\delta 1}$,...,$N_{\delta e}$ may optionally be provided for gathering the answers and merging the search results into a final result, in which case the dispatch node N_α, of course, is relieved of the acquisition task.

A first embodiment of the search engine according to the invention is shown in fig. 1, where the search nodes N_β are provided in columns or groups S provided and connected in parallel between the dispatch nodes N_α and the indexing nodes N_γ. The arrangement of the dispatch nodes N_α is shown schematically, but in practice they would be provided in a multilevel hierarchical arrangement.

A second embodiment of the search engine according to the invention is shown in fig. 2, where acquisition nodes N_δ are provided in a similar arrangement as that of the dispatch nodes, relieving the latter of the acquisition task.

It is to be understood that singular workstations may implement a specific type of nodes only, on alternatively more than one type of nodes. In other words, the different types of nodes may be distributed over the cluster of workstations. Hence the architecture shown in figs. 1 and 2 is implemented by the whole cluster, and these figures accordingly do neither show the workstations nor the interconnect bus.

The nodes shall now be discussed in more detail, starting with the search nodes which are central to the search engine according to the invention.

A search node N_β holds as mentioned a portion of the entire data set d_p. The search node has both a software search engine SW, and optionally a number of PMC modules M, as shown in fig. 3. The data set for a search node N_β is generated on an indexing node N_γ, of which more anon.

A search node may be equipped with a number x of PMC modules M for very fast searching, as shown in fig. 3. Each PMC module M has y groups G of z PMCs, as shown in figure 4, where each group G receives data from a single memory chip RAM. These A module M will typically be an individual circuit board. Each chip PMC is capable of processing q simultaneous queries, as shown schematically in fig. 5.

A pattern matching chip PMC can process a data volume of t_c bytes per second. Assuming that the memory modules are capable of delivering t_y bytes per second to the pattern matching chips PMCs, a PMC can search through the data volume of T_c bytes, $T_c = \min\{t_c, t_y\}t$, in the given time t.

As shown in figure 4, the pattern matching chips PMC are placed in modules M with y groups G of z chips PMC, where each group G receives data from a single memory chip RAM, and the size of the memory chip is T_c. The total amount of data this module can search through is $T_y = T_c$ with zq different queries.

When x modules M are provided in a search node N_β, these PMC modules M can search through an amount of data equal to $T_r = T_y \cdot x = \min\{ t_{c,} t_y\}txy$ – since no PMC modules search through the same data, the number of concurrent queries is still zq.

Thus the total query rate of the PMC modules in a search node can be expressed as

$$r_{HW} \frac{\frac{zq}{T_r}}{\min\{t_c, t_y\}\alpha y} \qquad (1)$$

where T_r denotes the total data volume on a node. The search node performance can now be calculated.

Given that the PMC modules M (or any hardware equivalent) has a query rate of r_{HW} and that the search software on a search node N_β has a query rate of r_{SW}, the total query rate r_Σ of a search node N_s can be expressed as

$$r_\Sigma = r_{t,HW}(1 - \varphi_{SW}) + r_{t,SW}\varphi_{SW} \qquad (2)$$

where φ_{SW} denotes the percentage of queries q that will be executed in software. The actual value of φ_{SW} is dynamically updated at runtime from a statistical model.

The dispatch nodes N_α receive all the queries, and resend them to all the search nodes N_β. The answers from the different search nodes N_β are merged and in case the dispatch nodes N_α functions as acquisition nodes, a complete answer is returned.

The indexing nodes N_γ collect documents and create prebuilt indexes for the search software on the different search nodes N_β. Hence the indexing nodes N_γ can be incorporated in the search nodes N_β with appropriate indexing software in the latter. The hardware is based on scanning through the entire collection of raw data, but some preprocessing and filtering of the raw data can be done in the indexing nodes N_γ.

Concerning the interconnect and data traffic, some general observations can be made based on the following considerations.

Different types of interconnect can be used for connecting the nodes. For a lower end system, a regular 100 Mbit Fast Ethernet will for instance handle the traffic.

Traffic on the interconnect between the nodes can be divided into two categories:

- Query traffic – traffic between dispatch nodes N_α and search nodes N_β. This traffic is present all the time when searching is performed. The query

traffic is characterized by low to medium data volumes and high frequency.

- Data traffic – traffic between indexing nodes N_γ and search nodes N_β. The data traffic is characterized by high data volumes and low frequency (typically one batch per day).

A typical query will transfer a query string from the dispatch node N_α to the search nodes N_β. Then the search nodes N_β will reply with a sequence of documents matching the query. Optionally N_α shall also be able to query the search node for the URL strings for the document, but this is considered immaterial in the present context.

The architecture of the search engine according to the invention can, based on the above considerations, now easily be scaled in two dimensions, viz. the data volume and the performance dimension respectively,

Data volume scaling is achieved by adding more data set partitions *d*, in other words more groups or columns S of search nodes N_β are added. Also the number of indexing nodes N_γ and dispatch nodes N_α can be increased as necessary in order to handle more data set partitions *d*.

Performance scaling can be achieved in the search engine architecture by replicating data set partitions *d* with a corresponding increase in the number of search nodes N_β, by increasing the number as illustrated in fig. 6. When using replication of data set partitions to scale the system performance, each search node N_β is part of the search node group S. Thus, the search nodes $N_{\beta 1},\ldots,N_{\beta s}$ are arranged into groups $S_{\beta 1},\ldots,S_{\beta v}$, where $v = \frac{s}{h_s}$, h_s denoting the scaling factor. The group $S_{\beta j}$ contains the search nodes $N_{\beta,j}, N_{\beta,j+1}$, $N_{\beta,j+2}$ and $N_{\beta,j+3}$ as rendered in fig. 8.

Scaling the data volume may cause the number of search nodes N_β receiving queries broadcast from a dispatch node N_α growing quite large. The architecture solves this problem by using several levels λ of dispatch nodes N_α – this is illustrated in fig. 7, which renders the arrangement of the dispatch nodes N_α as nodes in a portion of a binary data distribution tree. A binary data distribution tree easily allows for a linear scalability. Similar binary data distribution trees of the kind that already has been disclosed in

the applicant's above-mentioned international application PCT/NO99/00344, which discloses the configuration of an actual implementation of the pattern matching chip PMC. The number of dispatch nodes N_α in a regular binary tree is, of course $2^{\lambda-1}$ on each level λ, $\lambda = 1,2,3...$ A dispatch root node is on the first level, and up to and including a given level λ there is a total of $2^\lambda-1$ dispatch nodes in the tree. In case the dispatch nodes N_α are used also as acquisition nodes, i.e. for gathering the answers returned from the search nodes, the results of a search is merged in the dispatch nodes, the root dispatch node outputting the final answer to the query. However, there is nothing against that the search engine according to the invention is set up with a separate data gathering tree connected to the search nodes and comprising acquisition nodes N_δ gathering and outputting the final result of a query on the acquisition root node of the data gathering tree, i.e. the data acquisition node tree. The acquisition node tree could then be a mirror image of the dispatch node tree.

The schematic layout of a scalable search engine architecture according to the invention is shown in fig. 8, with the principle for two-dimensional scaling illustrated. It will be seen that the dispatch nodes N_α constitute a front end of the search engine and route the queries to the search node N_β and receives the search results back from the search nodes, wherein the actual search of the indexed data is performed. In case dedicated acquisition nodes N_δ are used, as shown in fig. 9, which otherwise is similar to fig. 8, the search results would, of course, be returned thereto. With the dispatch nodes N_α arranged in a tree configuration as shown in fig. 9, the acquisition node network as a back end of the search engine would form a mirror image of the dispatch node network. The indexing (spidering) nodes N_γ also constitute a back end of the search engine and collect data from e.g. the Internet and index the data to generate a searchable catalogue. By adding search nodes N_β or rather search nodes groups S horizontally the search engine scales linearly in the data volume, each additional search node or search node group containing different data. Typical capacity parameters for a search engine can as a non-limiting instance be given as follows. One search node N_β can typically handle 8 000 000 page views per day in a 5 000 000 documents catalogue. For a scalable search engine each search node N_β typically could hold 5 000 000 unique indexed documents, implying that 40 search nodes in one row are enough to maintain a 200 000 000 documents catalogue. Scaling

the performance, i.e. increasing the traffic capacity demands more rows of search nodes N_β with the same data to be added, such that the search nodes in a single column or group S contain identical data. A group or column S of 10 search nodes N_β hence will be able to handle 80 000 000 page views per day, with 40 columns handling a total of 3 200 000 000 page views per day.

An additional important benefit of a search engine according to the invention with the architecture scalable as herein disclosed, is that the query response time is essentially independent of the catalogue size, as each query is executed in parallel on all search nodes N_β and that the architecture is inherently fault-tolerant, such that faults in the individual nodes will not result in a system breakdown, only temporary reduce the performance until the fault is corrected.

Moreover, the in principle unlimited linear scalability of the data volume and the traffic volume which can be provided in a search engine according to the invention, contrasts sharply with prior art search engines, wherein the search cost typically increases exponentially with the data or traffic volume increase, and wherein the maximum capacity of the prior art search engines typically will be reached at low to moderate volumes. With the search engine according to the invention the cost will scale linearly with the increase in capacity at most, depending actually on whether the capacity increase is provided by adding SW search nodes only or also SW/HW search nodes. Finally the search engine according to the invention offers the advantage that each node in practice can be implemented with standard low cost commercially available PCs, but alternatively also with more expensive UNIX-based servers such as for instance the Sun or Alpha computers as currently available.

PATENT CLAIMS

1. A search engine with two-dimensional linearly scalable parallel architecture for searching a collection of text documents D, wherein the documents can be divided into a number of partitions $d_1, d_2, \ldots d_n$, wherein the collection of documents D is preprocessed in a text filtration system such that a preprocessed document collection D_p is obtained and corresponding preprocessed partitions $d_{p1}, d_{p2}, \ldots d_{pn}$, wherein an index I can be generated from the document collection D such that for each previous preprocessed partition $d_{p1}, d_{p2}, \ldots d_{pn}$ a corresponding index $i_1, i_2, \ldots i_n$ is obtained, wherein searching a partition d of the document collection D takes place with a partition-dependent data set $d_{p,k}$ comprising both the preprocessed partition d_{pk} and the corresponding index i_k, with $1 \leq k \leq n$, and wherein
the search engine comprises data processing units which form sets of nodes (N) connected in a network, characterized in that
the first set of nodes comprises a dispatch nodes ($N_{\alpha 1}, \ldots, N_{\alpha a}$), a second set of nodes comprises b search nodes ($N_{\beta 1}, \ldots, N_{\beta b}$), a third set of nodes comprises g indexing nodes ($N_{\gamma 1}, \ldots, N_{\gamma g}$), and an optional fourth set of nodes comprises e acquisition nodes ($N_{\delta 1}, \ldots N_{\delta e}$).
that the dispatch nodes (N_α) are connected in a multilevel configuration in the network,
that the search nodes (N_β) are grouped in v columns (S) which are connected in parallel in the network between the dispatch nodes (N_α) and an indexing node (N_γ), that the dispatch nodes (N_α) are adapted to process search queries and search answers, the search queries being dispatched further to all search nodes (N_β), and in case the acquisition nodes (N_δ) are not present, the search answers being returned to the dispatch nodes (N_α) and therein being combined to a final search result,
that the search nodes (N_β) each are adapted to contain search software,
that the indexing nodes (N_γ) are adapted for generally generating indexes i for the search software and optionally for generating partition-dependent data sets $d_{p,k}$ to search nodes (N_β) which comprise a search processor module,
that in case acquisition nodes (N_δ) are present, these are connected in a multilevel configuration in the network similar to that of the dispatch nodes (N_α), and adapted for gathering answers to search queries and outputting a final result thereof, thus relieving the dispatch nodes of this task,

that the number v of search node columns (S) scales with the number n of partitions d, thus implementing a scaling of the data volume, and that the number of search nodes (N_β) in each search node column (S_β) scales with an estimated or expected traffic load, thus implementing a scaling of performance, whereby in any case searching the document collection D takes place with a single search node (N_β) holding one of the partition-dependent data sets $d_{p,k}$ and all search nodes (N_β) in a search node column (S_β) containing identical partition-dependent data sets $d_{p,k}$.

2. A search engine according to claim 1, characterized in that the multilevel configuration of the dispatch nodes (N_α) and the optional acquisition nodes (N_δ) in the network are provided by hierarchical tree structures.

3. A search engine according to claim 2, characterized in that multilevel configuration of the optional acquisition nodes (N_δ) is a mirror image of the multilevel configuration of the dispatch nodes (N_α).

4. A search engine according to claim 2, characterized in that the hierarchical tree structures are binary tree structures.

5. A search engine according to claim 1, characterized in that each of the search nodes (N_β) comprises a search software module (SW).

6. A search engine according to claim 5, characterized in that at least some of the search nodes (N_β) comprises at least one dedicated search processor module (M), each dedicated search processor module (M) being realized with one or more dedicated search processor chips (PMC) which each is adapted for parallel handling of a number of q search queries.

7. A search engine according to claim 6, characterized in that the dedicated search processor chips (PMC) are provided in the search processor modules (M) in y processor groups (G), each with z search processor chips (PMC) and being connected with and adapted to receive data from a memory (RAM) assigned to the processor group (G).

8. A search engine according to claim 1, characterized in that an increase in the number of partitions *d* in the scaling of the data volume is accompanied by a corresponding increase in the number of dispatch nodes (N_α) and, in case, also in the number of acquisition nodes (N_δ), and optionally also by an increase in the number of index nodes (N_γ).

9. A search engine according to claim 1, characterized in that the separate node sets (N_α, N_β, N_γ, N_δ) each is implemented over one or more workstations connected in a data communications network.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	**ES**	Spain	**LS**	Lesotho	**SI**	Slovenia
AM	Armenia	**FI**	Finland	**LT**	Lithuania	**SK**	Slovakia
AT	Austria	**FR**	France	**LU**	Luxembourg	**SN**	Senegal
AU	Australia	**GA**	Gabon	**LV**	Latvia	**SZ**	Swaziland
AZ	Azerbaijan	**GB**	United Kingdom	**MC**	Monaco	**TD**	Chad
BA	Bosnia and Herzegovina	**GE**	Georgia	**MD**	Republic of Moldova	**TG**	Togo
BB	Barbados	**GH**	Ghana	**MG**	Madagascar	**TJ**	Tajikistan
BE	Belgium	**GN**	Guinea	**MK**	The former Yugoslav	**TM**	Turkmenistan
BF	Burkina Faso	**GR**	Greece		Republic of Macedonia	**TR**	Turkey
BG	Bulgaria	**HU**	Hungary	**ML**	Mali	**TT**	Trinidad and Tobago
BJ	Benin	**IE**	Ireland	**MN**	Mongolia	**UA**	Ukraine
BR	Brazil	**IL**	Israel	**MR**	Mauritania	**UG**	Uganda
BY	Belarus	**IS**	Iceland	**MW**	Malawi	**US**	United States of America
CA	Canada	**IT**	Italy	**MX**	Mexico	**UZ**	Uzbekistan
CF	Central African Republic	**JP**	Japan	**NE**	Niger	**VN**	Viet Nam
CG	Congo	**KE**	Kenya	**NL**	Netherlands	**YU**	Yugoslavia
CH	Switzerland	**KG**	Kyrgyzstan	**NO**	Norway	**ZW**	Zimbabwe
CI	Côte d'Ivoire	**KP**	Democratic People's	**NZ**	New Zealand		
CM	Cameroon		Republic of Korea	**PL**	Poland		
CN	China	**KR**	Republic of Korea	**PT**	Portugal		
CU	Cuba	**KZ**	Kazakstan	**RO**	Romania		
CZ	Czech Republic	**LC**	Saint Lucia	**RU**	Russian Federation		
DE	Germany	**LI**	Liechtenstein	**SD**	Sudan		
DK	Denmark	**LK**	Sri Lanka	**SE**	Sweden		
EE	Estonia	**LR**	Liberia	**SG**	Singapore		

A search engine with two-dimensional linearly scalable parallel architecture.

The present invention concerns a search engine with two-dimensional linearly scalable parallel architecture for searching a collection of text documents D, wherein the documents can be divided into a number of partitions $d_1, d_2, \ldots d_n$, wherein the collection of documents D is preprocessed in a text filtration system such that a preprocessed document collection D_p is obtained and corresponding preprocessed partitions $d_{p1}, d_{p2}, \ldots d_{pn}$, wherein an index I can be generated from the document collection D such that for each previous preprocessed partition $d_{p1}, d_{p2}, \ldots d_{pn}$ a corresponding index $i_1, i_2, \ldots i_n$ is obtained, wherein searching a partition d of the document collection D takes place with a partition-dependent data set $d_{p,k}$ comprising both the preprocessed partition d_{pk} and the corresponding index i_k, with $1 \leq k \leq n$, and wherein the search engine comprises data processing units which form sets of nodes connected in a network.

Most prior art search engines work with large data set and employ powerful computers to perform the search. However, searching is a partitionable data processing problem, and this fact can be used to partition a search problem into a large number of specific queries and let each query be processed simultaneously on a commensurate number of processors connected in parallel in a network. Particularly searching can be regarded as a binary partitionable data processing problem, and hence a binary tree network is used for establishing a multiprocessor architecture such as disclosed for instance in US patent No. 4860201 (Stolfo & al.) and international patent application PCT/NO99/00308 which belongs to the applicant and hereby is incorporated by reference. The present applicant has developed proprietary technologies for searching within regular text documents. These technologies are i.a. based on a search system and a method for searching as described in international patent application PCT/NO99/00233 which belongs to the applicant and hereby is incorporated by reference. The search system is based on efficient core search algorithms which may be used in the search engine according to the invention.

However, it has become increasingly important to cater for a growing number of documents to be searched and also to be able to handle an increased traffic load, i.e. the number of queries per second which shall be processed by the

search system. This, apart from the ability to handle a large number of queries simultaneously on processor level, implies that a search engine should be implemented with an architecture that allows for preferably linear scalability in two dimensions, viz. both with regard to the data volume and the performance, i.e. the ability to handle a very large number of queries per second. Considering the development of the World Wide Web, a scalability problem in the search engine architecture will be extremely important as there presently is an enormous growth rate in both the number of documents and the number of users on the Internet.

Prior art search engine solutions for Internet are able to scale to a certain level, but almost always this is achieved in a manner that requires a high cost increase of the search engine system relative to the growth in data volume or data traffic. Very often the system costs scale as the square of the data volume or the traffic, a doubling of the data volume thus leading to quadrupled system costs. Furthermore all the major Internet search engines presently are based on very expensive server technology, often coupled with brute computing force-approaches and accompanied with disadvantages such as slow server turnaround, requirements for special hardware to provide fault tolerance etc. The system costs can e.g. be measured as the amount of hardware required to implement a search engine solution or the actual aggregated price of the system.

It is thus a main object of the present invention to provide a search engine with a multilevel data and functional parallelism, such that large volumes of data can be searched efficiently and very fast by a large number of users simultaneously.

Particularly it is a further object of the invention to provide a parallel architecture for implementing a search engine with a multilevel data and functional parallelism.

Yet a further object of the present invention is to provide a parallel architecture which is linearly scalable in two dimensions, i.e. with regard to both data volume and performance, that is the query rate.

The above-mentioned objects and further features and advantages are provided with a search engine according to the invention which is characterized in that the first set of nodes comprises *a* dispatch nodes, a

second set of nodes comprises b search nodes, a third set of nodes comprises g indexing nodes, and an optional fourth set of nodes comprises e acquisition nodes, that the dispatch nodes are connected in a multilevel configuration in the network, that the search nodes are grouped in columns which are connected in parallel in the network between the dispatch nodes and an indexing node, that the dispatch nodes are adapted to process search queries and search answers, the search queries being dispatched further to all search nodes and in case the acquisition nodes are not present, the search answers being returned to the dispatch nodes and therein being combined to a final search result, that the search nodes each are adapted to contain search software, that at least some of the search nodes additionally comprise at least one search processor module, that the indexing nodes are adapted for generally generating indexes i for the search software and optionally for generating partition-dependent data sets $d_{p,k}$ to search nodes which comprise a search processor module, that in case acquisition nodes are present, these are connected in a multilevel configuration in the network similar to that of the dispatch node, and adapted for gathering answers to search queries and outputting a final result thereof, thus relieving the dispatch nodes of this task, and that the two-dimensional linear scaling respectively takes place by scaling of the data volume through an increase in the number of partitions d and scaling of performance through replication of one or more partitions d.

According to the invention are advantageously the multilevel configuration of the dispatch nodes and the optional acquisition nodes network provided by hierarchical tree structures, and the multilevel configuration of the optional acquisition nodes is then preferably a mirror image of the multilevel configuration of the dispatch nodes, the hierarchical tree structures preferably being binary tree structures.

According to the invention each of the search nodes advantageously comprises a search software module.

Further, according to the invention at least some of the search nodes comprises at least one dedicated search processor module, each dedicated search processor module being realized with one or more dedicated processor chips, which is adapted for parallel handling of a number of search queries. In this connection it is preferred that the dedicated search processor chips are provided in the search processor modules in y processor groups, each with z

search processor chips and being connected with and adapted to receive data from a memory assigned to the processor group.

According to the invention the increase in the number of partitions in the scaling of the data volume is advantageously implemented by an increase in the number of search node groups or columns. In this connection the increase in the number of partitions can preferably be accompanied by a corresponding increase in the number of dispatch nodes and. in case, also in the number of acquisition nodes, and optionally also by an increase in the number of index nodes.

According to the invention the replication of one or more partition in the scaling of performance is advantageously implemented by an increase in the number of search nodes in each group or column.

Finally it is according to the invention advantageous that the separate node sets each is implemented over one or more workstations connected in a data communications network.

The search engine according to the invention shall now be described in terms of non-limiting exemplary embodiments and with reference to the accompanying drawings, in which

fig. 1 shows a general overview of the architecture in a first embodiment of the search engine according to the invention.

fig. 2 a general overview of the architecture in a second embodiment of the search engine according to the invention,

fig. 3 schematically a search node with search software and dedicated hardware search chips,

fig. 4 the layout of a single module of search chips,

fig. 5 schematically the query handling in a dedicated search chip,

fig. 6 schematically the principle of performance scaling.

fig. 7 a multilevel dispatch node configuration based on a binary tree structure.

fig. 8 an overview of the architecture of a first embodiment of the search engine according to the invention. with an indication of the principle for two-dimensional scaling, and

fig. 9 an overview of the architecture of a second embodiment of the search engine according to the invention. with an indication of the principle for two-dimensional scaling.

Searching a large collection of independent documents is a highly parallel task. The search engine according to the invention employs parallelism on different levels as shall be discussed in the following.

The search engine according to the present invention searches a document collection of documents D. The documents can be divided into n partitions $d_1, d_2, \ldots, d_n$. Each document collection D. or partition d of a document collection can be preprocessed for use in a hardware text filtering system, for instance implemented by dedicated hardware like the applicant's so-called Pattern Matching Chip (PMC) which is disclosed in the applicant's international patent application No. PCT/NO99/00344 which hereby is incorporated by reference. The preprocessed document collection is denoted D_p and the corresponding preprocessed document collection partitions d_p are denoted $d_{p1}, d_{p2}, \ldots, d_{pn}$.

Software-based search systems require an index generated from the document collection. The index is denoted I and the indexes corresponding to the document collection partitions $d_{p1}, d_{p2}, \ldots d_{pn}$, are denoted $i_1, i_2, \ldots i_n$.

The data set needed for searching a partition d of the document collection D, is called the partition-required or partition-dependent data set. In a software only system (SW system) the data set is the index i_k, while in the systems with hardware (SW/HW systems). the data set also includes the preprocessed document collection partition $d_{p.k}$ with the corresponding index i_k, where $1 \leq k \leq n$.

The essentially software-implemented partitioning and preprocessing operations can be rendered schematically as $*(D) \rightarrow *(d_1, \ldots, d_n) \rightarrow *(d_{p1}, \ldots, d_{pn}) \rightarrow d_{p.k}$, where $*(D)$ denotes a partitioning operation on the input D. $*(d_1, \ldots, d_n)$ a filtering operation. e.g. indexing, on $d_1, \ldots, d_n$ and $d_{p.k}$ of course is the partition-dependent data set. which in an SW system only shall be the index i_k. and with $1 \leq k \leq n$.

A search engine is implemented on a cluster of workstations that are connected using a high performance interconnect bus. The not shown workstations then constitute the server of the search system (search server). The workstations implement nodes of the search server. The nodes perform different tasks and are according to the invention implemented as set out below.

The nodes can be regarded as virtual nodes distributed among the workstations, but in a SW/HW search engine the dedicated search processing hardware must be physically present in some workstations in order to support the hardware-based search nodes. The search node software may then still be distributed. Also, some search nodes in SW/HW search engine may comprise software only, and optionally be distributed over more than one workstation.

- Dispatch nodes $N_{\alpha 1},\ldots N_{\alpha a}$ handle incoming queries and send queries along to all the search nodes. The dispatch nodes can also be configured as acquisition nodes, gathering the answers, i.e. the search results to the queries. Upon receiving the answers, the dispatch nodes in the acquisition mode merge the search results into a final result.
- Search nodes $N_{\beta 1},\ldots,N_{\beta b}$ hold a portion of the entire data set d_p. A search node comprises both the dedicated search software as well as a number of the above-mentioned PMC modules for hardware searching.
- Indexing nodes $N_{\gamma 1},\ldots,N_{\gamma g}$ are responsible for generating indexes for the dedicated search software on a number of search nodes. For the PMC modules, the indexing node N_γ also filters and preprocesses the raw data.
- Acquisition nodes $N_{\delta 1},\ldots,N_{\delta e}$ may optionally be provided for gathering the answers and merging the search results into a final result, in which case the dispatch node N_α, of course, is relieved of the acquisition task.

A first embodiment of the search engine according to the invention is shown in fig. 1, where the search nodes N_β are provided in columns or groups S provided and connected in parallel between the dispatch nodes N_α and the indexing nodes N_γ. The arrangement of the dispatch nodes N_α is shown schematically, but in practice they would be provided in a multilevel hierarchical arrangement.

A second embodiment of the search engine according to the invention is shown in fig. 2. where acquisition nodes N_δ are provided in a similar arrangement as that of the dispatch nodes. relieving the latter of the acquisition task.

It is to be understood that singular workstations may implement a specific type of nodes only, on alternatively more than one type of nodes. In other words, the different types of nodes may be distributed over the cluster of workstations. Hence the architecture shown in figs. 1 and 2 is implemented by the whole cluster, and these figures accordingly do neither show the workstations nor the interconnect bus.

The nodes shall now be discussed in more detail. starting with the search nodes which are central to the search engine according to the invention.

A search node N_β holds as mentioned a portion of the entire data set d_p. The search node has both a software search engine SW, and optionally a number of PMC modules M, as shown in fig. 3. The data set for a search node N_β is generated on an indexing node N_γ, of which more anon.

A search node may be equipped with a number x of PMC modules M for very fast searching, as shown in fig. 3. Each PMC module M has y groups G of z PMCs, as shown in figure 4, where each group G receives data from a single memory chip RAM. These A module M will typically be an individual circuit board. Each chip PMC is capable of processing q simultaneous queries. as shown schematically in fig. 5.

A pattern matching chip PMC can process a data volume of t_c bytes per second. Assuming that the memory modules are capable of delivering t_y bytes per second to the pattern matching chips PMCs, a PMC can search through the data volume of T_c bytes, $T_c = \min\{t_c, t_y\}t$, in the given time t.

As shown in figure 4. the pattern matching chips PMC are placed in modules M with y groups G of z chips PMC. where each group G receives data from a single memory chip RAM, and the size of the memory chip is T_c. The total amount of data this module can search through is $T_y = T_c$ with zq different queries.

When x modules M are provided in a search node N_β, these PMC modules M can search through an amount of data equal to $T_r = T_y \cdot x = \min\{ t_c, t_y\}txy$ –

since no PMC modules search through the same data, the number of concurrent queries is still zq.

Thus the total query rate of the PMC modules in a search node can be expressed as

$$r_{HW} \frac{zq}{\frac{T_r}{\min\{t_c, t_y\} txy}} \qquad (1)$$

where T_r denotes the total data volume on a node. The search node performance can now be calculated.

Given that the PMC modules M (or any hardware equivalent) has a query rate of r_{HW} and that the search software on a search node N_β has a query rate of r_{SW}, the total query rate r_Σ of a search node N_s can be expressed as

$$r_\Sigma = r_{t,HW}(1 - \varphi_{SW}) + r_{t,SW}\varphi_{SW} \qquad (2)$$

where φ_{SW} denotes the percentage of queries q that will be executed in software. The actual value of φ_{SW} is dynamically updated at runtime from a statistical model.

The dispatch nodes N_α receive all the queries, and resend them to all the search nodes N_β. The answers from the different search nodes N_β are merged and in case the dispatch nodes N_α functions as acquisition nodes, a complete answer is returned.

The indexing nodes N_γ collect documents and create prebuilt indexes for the search software on the different search nodes N_β. Hence the indexing nodes N_γ can be incorporated in the search nodes N_β with appropriate indexing software in the latter. The hardware is based on scanning through the entire collection of raw data, but some preprocessing and filtering of the raw data can be done in the indexing nodes N_γ.

Concerning the interconnect and data traffic, some general observations can be made based on the following considerations.

Different types of interconnect can be used for connecting the nodes. For a lower end system, a regular 100 Mbit Fast Ethernet will for instance handle the traffic.

Traffic on the interconnect between the nodes can be divided into two categories:

- Query traffic – traffic between dispatch nodes N_α and search nodes N_β. This traffic is present all the time when searching is performed. The query traffic is characterized by low to medium data volumes and high frequency.
- Data traffic – traffic between indexing nodes N_γ and search nodes N_β. The data traffic is characterized by high data volumes and low frequency (typically one batch per day).

A typical query will transfer a query string from the dispatch node N_α to the search nodes N_β. Then the search nodes N_β will reply with a sequence of documents matching the query. Optionally N_α shall also be able to query the search node for the URL strings for the document, but this is considered immaterial in the present context.

The architecture of the search engine according to the invention can, based on the above considerations, now easily be scaled in two dimensions, viz. the data volume and the performance dimension respectively,

Data volume scaling is achieved by adding more data set partitions *d*, in other words more groups or columns S of search nodes N_β are added. Also the number of indexing nodes N_γ and dispatch nodes N_α can be increased as necessary in order to handle more data set partitions *d*.

Performance scaling can be achieved in the search engine architecture by replicating data set partitions *d* with a corresponding increase in the number of search nodes N_β, by increasing the number as illustrated in fig. 6. When using replication of data set partitions to scale the system performance, each search node N_β is part of the search node group S. Thus, the search nodes $N_{\beta 1},\ldots,N_{\beta s}$ are arranged into groups $S_{\beta 1},\ldots,S_{\beta v}$, where $v=\frac{s}{h_s}$. h_s denoting the scaling factor. The group $S_{\beta j}$ contains the search nodes $N_{\beta,j}.N_{\beta,j+1}$, $N_{\beta,j+2}$ and $N_{\beta,j+3}$ as rendered in fig. 8.

Scaling the data volume may cause the number of search nodes N_β receiving queries broadcast from a dispatch node N_α growing quite large. The architecture solves this problem by using several levels λ of dispatch nodes

N_α – this is illustrated in fig. 7, which renders the arrangement of the dispatch nodes N_α as nodes in a portion of a binary data distribution tree. A binary data distribution tree easily allows for a linear scalability. Similar binary data distribution trees of the kind that already has been disclosed in the applicant's above-mentioned international application PCT/NO99/00344, which discloses the configuration of an actual implementation of the pattern matching chip PMC. The number of dispatch nodes N_α in a regular binary tree is, of course $2^{\lambda-1}$ on each level λ, $\lambda = 1,2,3...$ A dispatch root node is on the first level, and up to and including a given level λ there is a total of $2^\lambda-1$ dispatch nodes in the tree. In case the dispatch nodes N_α are used also as acquisition nodes, i.e. for gathering the answers returned from the search nodes, the results of a search is merged in the dispatch nodes, the root dispatch node outputting the final answer to the query. However, there is nothing against that the search engine according to the invention is set up with a separate data gathering tree connected to the search nodes and comprising acquisition nodes N_δ gathering and outputting the final result of a query on the acquisition root node of the data gathering tree, i.e. the data acquisition node tree. The acquisition node tree could then be a mirror image of the dispatch node tree.

The schematic layout of a scalable search engine architecture according to the invention is shown in fig. 8, with the principle for two-dimensional scaling illustrated. It will be seen that the dispatch nodes N_α constitute a front end of the search engine and route the queries to the search node N_β and receives the search results back from the search nodes, wherein the actual search of the indexed data is performed. In case dedicated acquisition nodes N_δ are used, as shown in fig. 9, which otherwise is similar to fig. 8, the search results would, of course, be returned thereto. With the dispatch nodes N_α arranged in a tree configuration as shown in fig. 9, the acquisition node network as a back end of the search engine would form a mirror image of the dispatch node network. The indexing (spidering) nodes N_γ also constitute a back end of the search engine and collect data from e.g. the Internet and index the data to generate a searchable catalogue. By adding search nodes N_β or rather search nodes groups S horizontally the search engine scales linearly in the data volume, each additional search node or search node group containing different data. Typical capacity parameters for a search engine can as a non-limiting instance be given as follows. One search node N_β can

typically handle 8 000 000 page views per day in a 5 000 000 documents catalogue. For a scalable search engine each search node N_β typically could hold 5 000 000 unique indexed documents, implying that 40 search nodes in one row are enough to maintain a 200 000 000 documents catalogue. Scaling the performance, i.e. increasing the traffic capacity demands more rows of search nodes N_β with the same data to be added, such that the search nodes in a single column or group S contain identical data. A group or column S of 10 search nodes N_β hence will be able to handle 80 000 000 page views per day, with 40 columns handling a total of 3 200 000 000 page views per day.

An additional important benefit of a search engine according to the invention with the architecture scalable as herein disclosed, is that the query response time is essentially independent of the catalogue size, as each query is executed in parallel on all search nodes N_β and that the architecture is inherently fault tolerant, such that faults in the individual nodes will not result in a system breakdown, only temporary reduce the performance until the fault is corrected.

Moreover, the in principle unlimited linear scalability of the data volume and the traffic volume which can be provided in a search engine according to the invention, contrasts sharply with prior art search engines, wherein the search cost typically increases exponentially with the data or traffic volume increase, and wherein the maximum capacity of the prior art search engines typically will be reached at low to moderate volumes. With the search engine according to the invention the cost will scale linearly with the increase in capacity at most, depending actually on whether the capacity increase is provided by adding SW search nodes only or also SW/HW search nodes. Finally the search engine according to the invention offers the advantage that each node in practice can be implemented with standard low cost commercially available PCs, but alternatively also with more expensive UNIX-based servers such as for instance the Sun or Alpha computers as currently available.

PATENT CLAIMS

1. A search engine with two-dimensional linearly scalable parallel architecture for searching a collection of text documents D, wherein the documents can be divided into a number of partitions $d_1, d_2, \ldots d_n$, wherein the collection of documents D is preprocessed in a text filtration system such that a preprocessed document collection D_p is obtained and corresponding preprocessed partitions $d_{p1}, d_{p2}, \ldots d_{pn}$, wherein an index I can be generated from the document collection D such that for each previous preprocessed partition $d_{p1}, d_{p2}, \ldots d_{pn}$ a corresponding index $i_1, i_2, \ldots i_n$ is obtained, wherein searching a partition d of the document collection D takes place with a partition-dependent data set $d_{p,k}$ comprising both the preprocessed partition d_{pk} and the corresponding index i_k, with $1 \leq k \leq n$, and wherein the search engine comprises data processing units which form sets of nodes (N) connected in a network, characterized in that the first set of nodes comprises a dispatch nodes ($N_{\alpha 1}, \ldots, N_{\alpha a}$), a second set of nodes comprises b search nodes ($N_{\beta 1}, \ldots, N_{\beta b}$), a third set of nodes comprises g indexing nodes ($N_{\gamma 1}, \ldots, N_{\gamma g}$), and an optional fourth set of nodes comprises e acquisition nodes ($N_{\delta 1}, \ldots N_{\delta e}$).
that the dispatch nodes (N_α) are connected in a multilevel configuration in the network,
that the search nodes (N_β) are grouped in columns (S) which are connected in parallel in the network between the dispatch nodes (N_α) and an indexing node (N_γ), that the dispatch nodes (N_α) are adapted to process search queries and search answers, the search queries being dispatched further to all search nodes (N_β), and in case the acquisition nodes (N_δ) are not present, the search answers being returned to the dispatch nodes (N_α) and therein being combined to a final search result,
that the search nodes (N_β) each are adapted to contain search software,
that at least some of the search nodes (N_β) additionally comprise at least one search processor module (M),
that the indexing nodes (N_γ) are adapted for generally generating indexes i for the search software and optionally for generating partition-dependent data sets $d_{p,k}$ to search nodes (N_β) which comprise a search processor module,
that in case acquisition nodes (N_δ) are present, these are connected in a multilevel configuration in the network similar to that of the dispatch nodes

(N_α), and adapted for gathering answers to search queries and outputting a final result thereof, thus relieving the dispatch nodes of this task, and that the two-dimensional linear scaling respectively takes place by scaling of the data volume through an increase in the number of partitions d and scaling of performance through replication of one or more partitions d.

2. A search engine according to claim 1, characterized in that the multilevel configuration of the dispatch nodes (N_α) and the optional acquisition nodes (N_δ) in the network are provided by hierarchical tree structures.

3. A search engine according to claim 2, characterized in that multilevel configuration of the optional acquisition nodes (N_δ) is a mirror image of the multilevel configuration of the dispatch nodes (N_α).

4. A search engine according to claim 2, characterized in that the hierarchical tree structures are binary tree structures.

5. A search engine according to claim 1, characterized in that each of the search nodes (N_β) comprises a search software module (SW).

6. A search engine according to claim 5, characterized in that at least some of the search nodes (N_β) comprises at least one dedicated search processor module (M), each dedicated search processor module (M) being realized with one or more dedicated search processor chips (PMC) which each is adapted for parallel handling of a number of q search queries.

7. A search engine according to claim 6, characterized in that the dedicated search processor chips (PMC) are provided in the search processor modules (M) in y processor groups (G), each with z search processor chips (PMC) and being connected with and adapted to receive data from a memory (RAM) assigned to the processor group (G).

8. A search engine according to claim 1, characterized in that the increase in the number of partitions d in the scaling of the data volume is implemented by an increase in the number of search node groups or columns (S).

9. A search engine according to claim 8.
characterized in that an increase in the number of partitions *d* in the scaling of the data volume is accompanied by a corresponding increase in the number of dispatch nodes (N_α) and. in case. also in the number of acquisition nodes (N_δ), and optionally also by an increase in the number of index nodes (N_γ).

10. A search engine according to claim 1,
characterized in that the replication of one or more partitions *d* in the scaling of the performance is implemented by an increase in the number of search nodes (N_β) in each group or column (S).

11. A search engine according to claim 1,
characterized in that the separate node sets (N_α, N_β, N_γ, N_δ) each is implemented over one or more workstations connected in a data communications network.

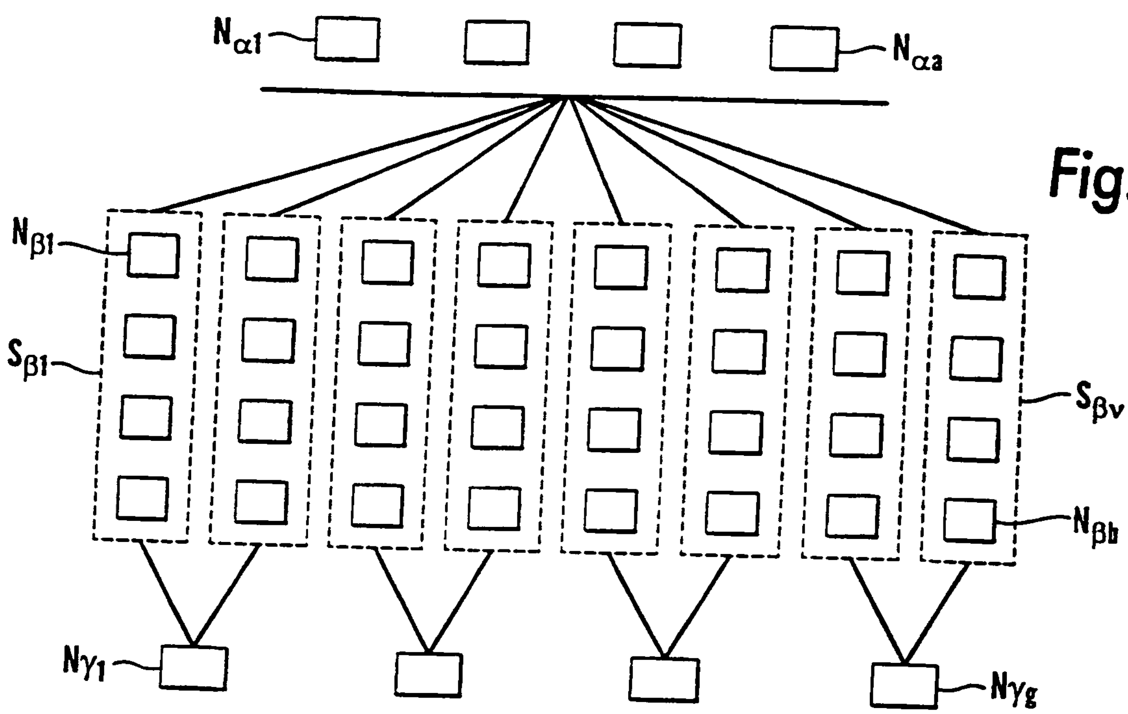
$N_{\alpha 1}$
$N_{\alpha a}$
$N_{\beta 1}$
$S_{\beta 1}$
$S_{\beta \nu}$
$N_{\beta b}$
$N_{\gamma 1}$
$N_{\gamma g}$

Fig. 1

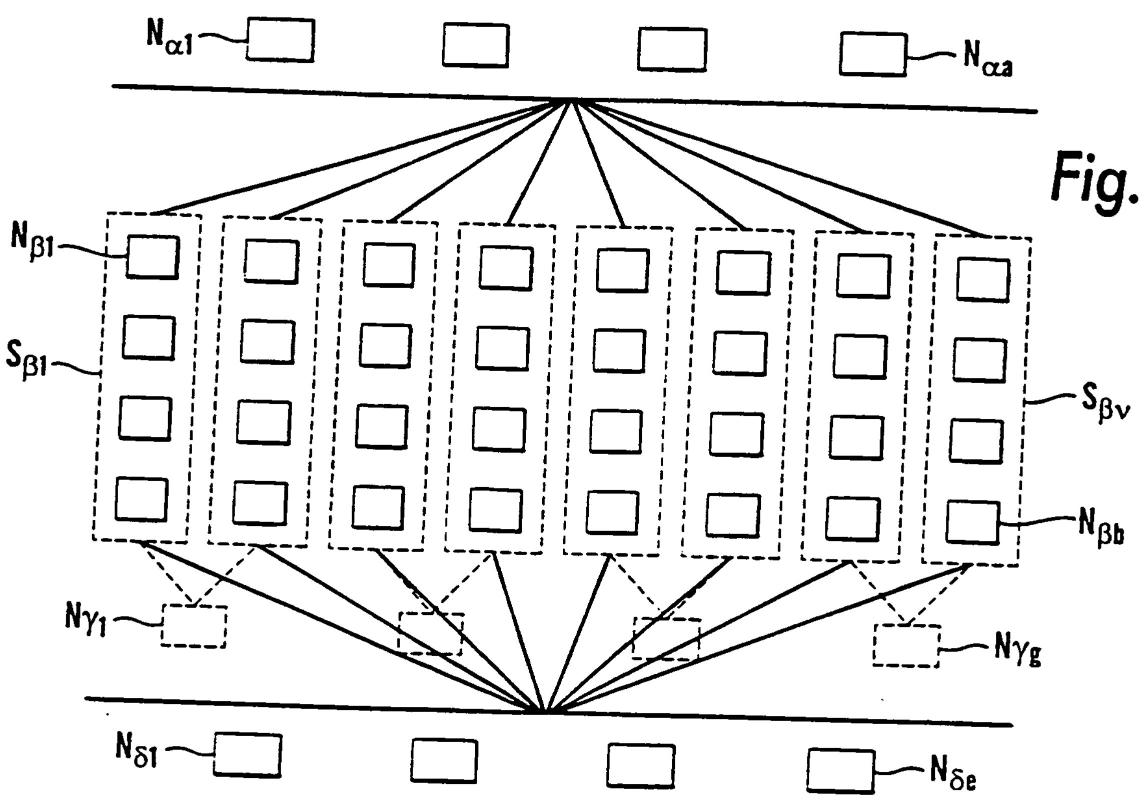
$N_{\alpha 1}$
$N_{\alpha a}$
$N_{\beta 1}$
$S_{\beta 1}$
$S_{\beta \nu}$
$N_{\beta b}$
$N_{\gamma 1}$
$N_{\gamma g}$
$N_{\delta 1}$
$N_{\delta e}$

Fig. 2

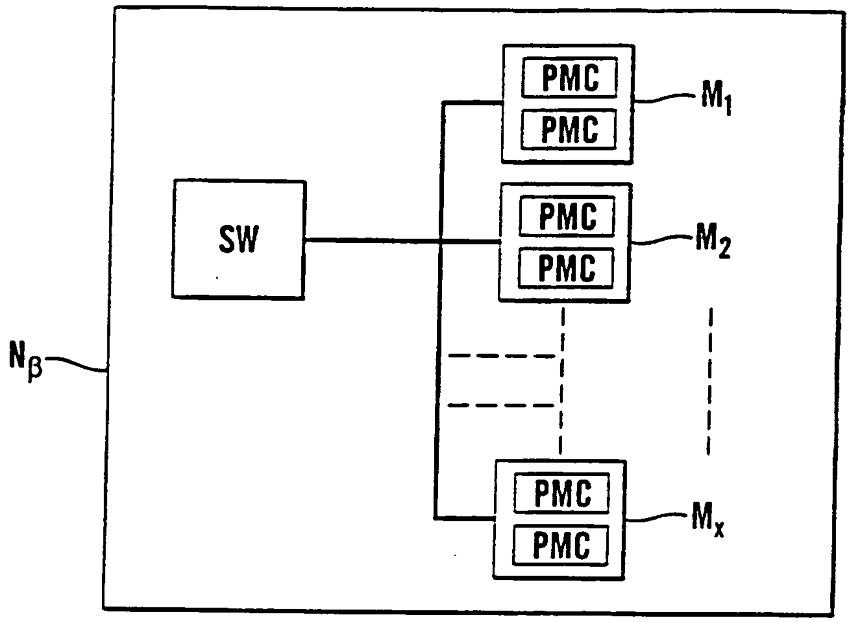
PMC
PMC
M_1
SW
PMC
PMC
M_2
N_β
PMC
PMC
M_x

Fig.3

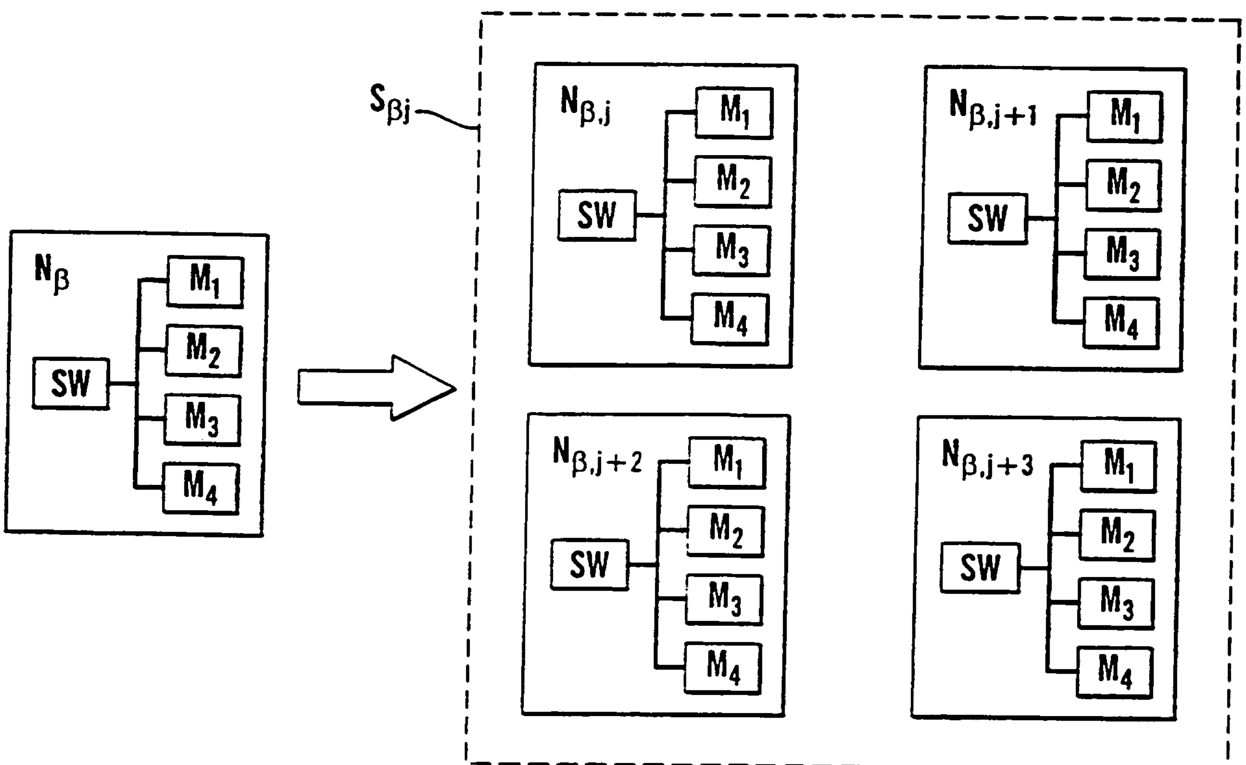
$S_{\beta j}$
$N_{\beta,j}$
SW
M_1
M_2
M_3
M_4
$N_{\beta,j+1}$
SW
M_1
M_2
M_3
M_4
N_β
SW
M_1
M_2
M_3
M_4
$N_{\beta,j+2}$
SW
M_1
M_2
M_3
M_4
$N_{\beta,j+3}$
SW
M_1
M_2
M_3
M_4

Fig.6

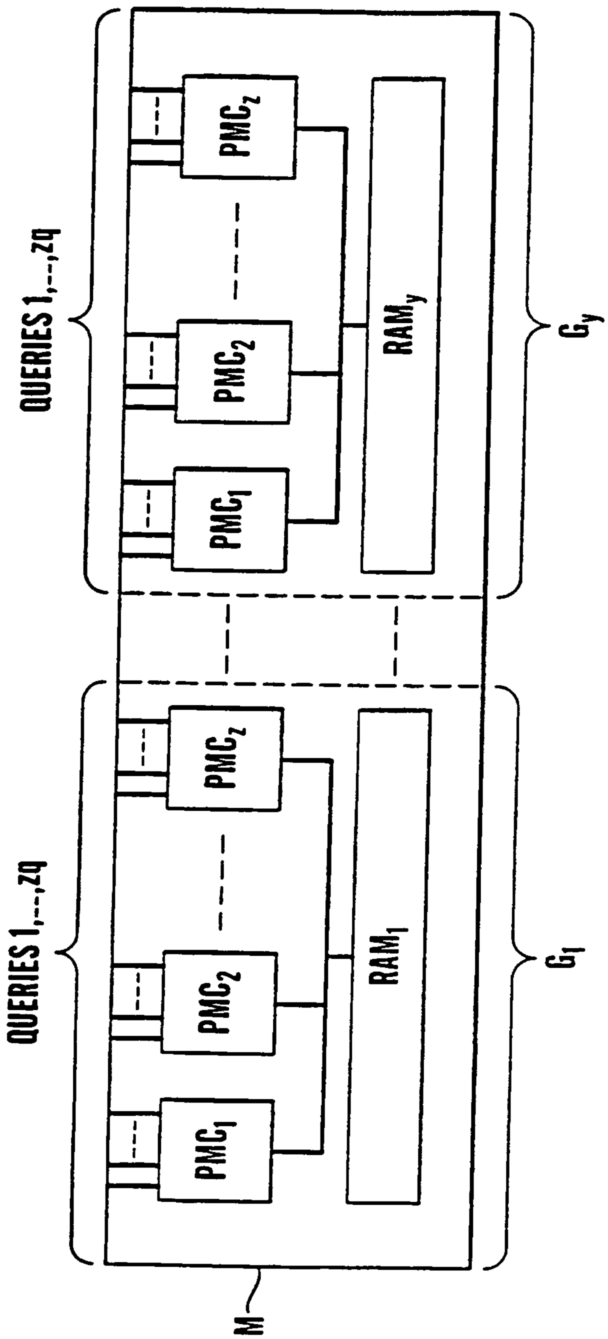
QUERIES 1,...,zq
PMC1
PMC2
PMCz
RAM1
G1
QUERIES 1,...,zq
PMC1
PMC2
PMCz
RAMy
Gy
M

Fig.4

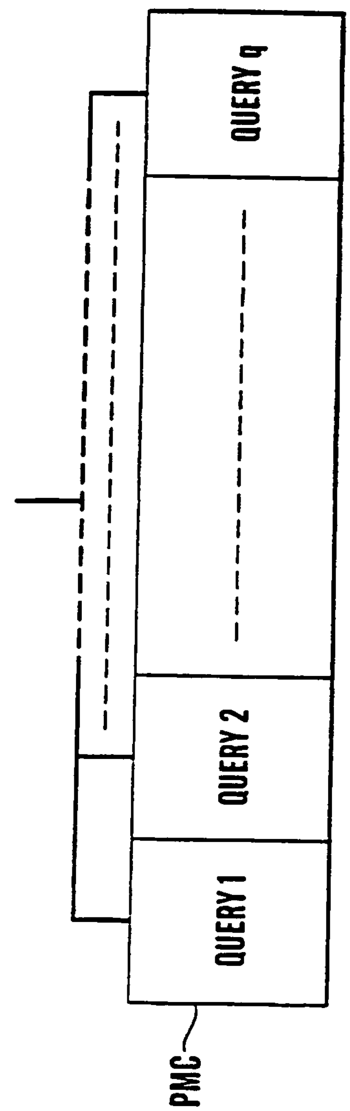
QUERY 1
QUERY 2
QUERY q
PMC

Fig.5

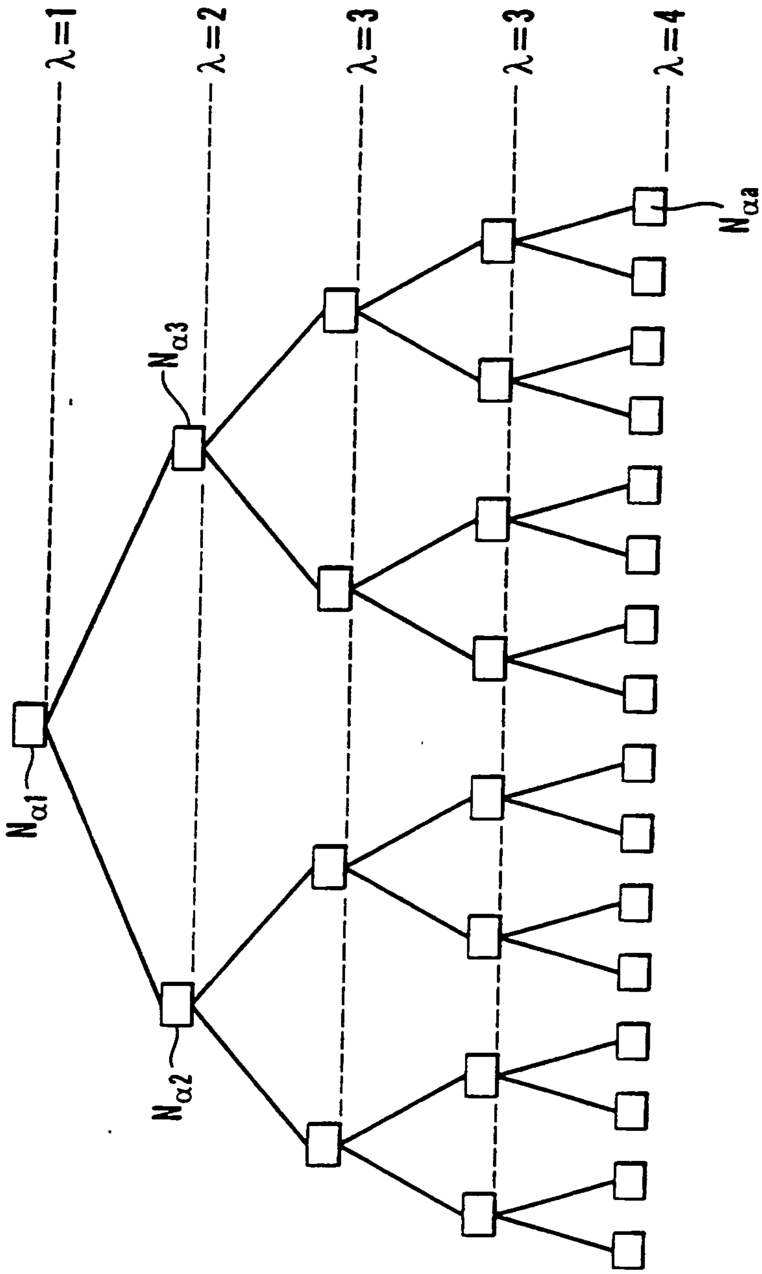
$\lambda=1$
$\lambda=2$
$\lambda=3$
$\lambda=3$
$\lambda=4$
$N_{\alpha 1}$
$N_{\alpha 2}$
$N_{\alpha 3}$
$N_{\alpha a}$

Fig.7

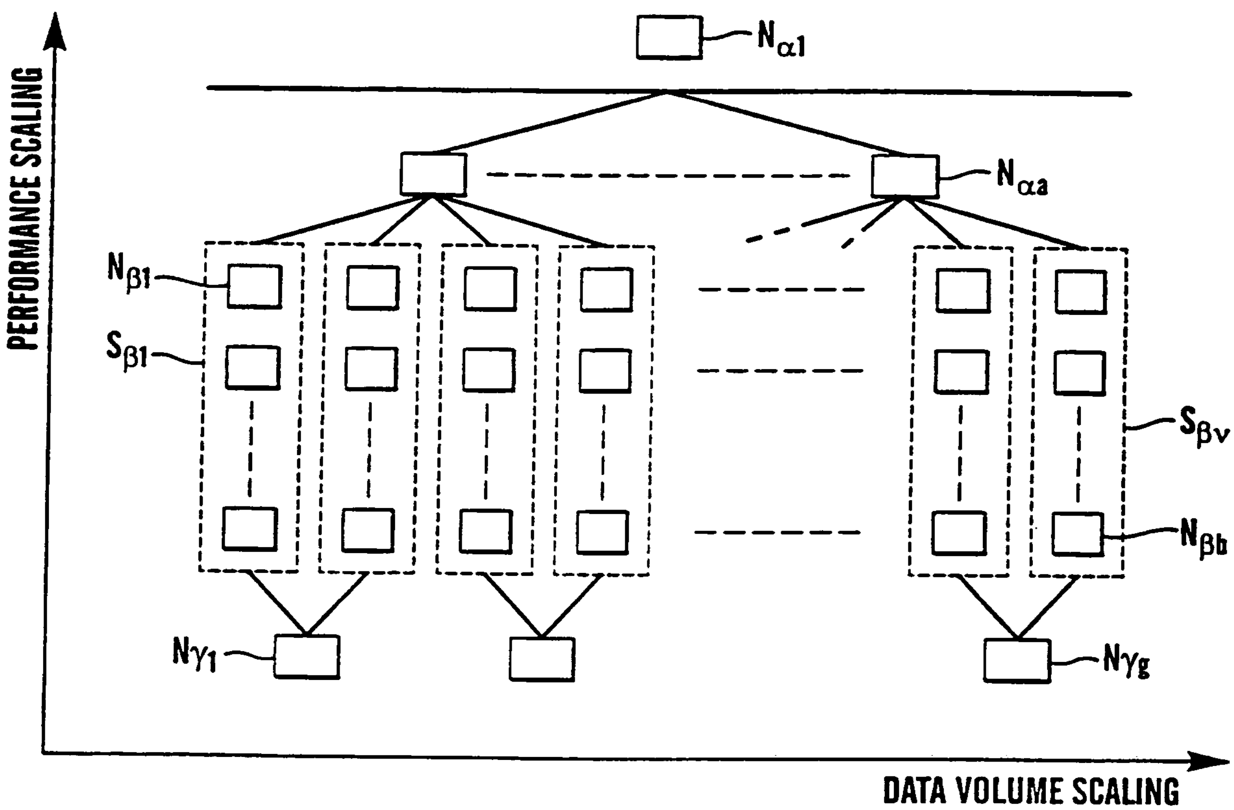
PERFORMANCE SCALING
$N_{\alpha 1}$
$N_{\alpha a}$
$N_{\beta 1}$
$S_{\beta 1}$
$S_{\beta v}$
$N_{\beta b}$
$N\gamma_1$
$N\gamma_g$
DATA VOLUME SCALING

Fig.8

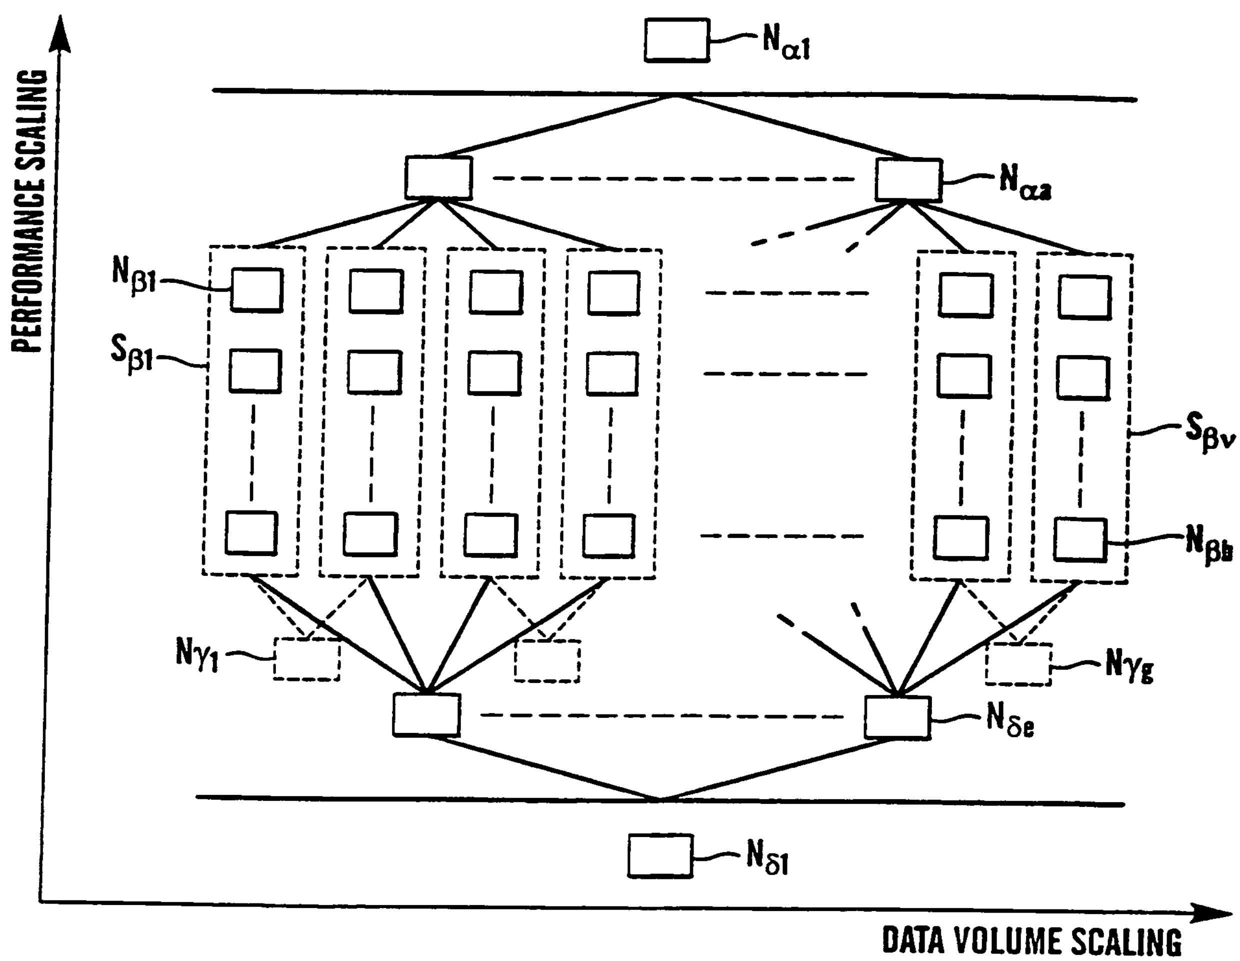
$N_{\alpha 1}$
$N_{\alpha a}$
$N_{\beta 1}$
$S_{\beta 1}$
$S_{\beta v}$
$N_{\beta b}$
$N_{\gamma 1}$
$N_{\gamma g}$
$N_{\delta e}$
$N_{\delta 1}$
PERFORMANCE SCALING
DATA VOLUME SCALING

Fig.9

09/743268

PCT

REQUEST

The undersigned requests that the present international application be processed according to the Patent Cooperation Treaty.

For receiving Office use only

International Application No.

International Filing Date

Name of receiving Office and "PCT International Application"

Applicant's or agent's file reference *(if desired) (12 characters maximum)* Opti44Fast7

Box No. I TITLE OF INVENTION

A search engine with two-dimensional linearly scalable parallel architecture

Box No. II APPLICANT

Name and address: *(Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)*

FAST SEARCH & TRANSFER ASA
P.O.Box 1677 Vika
N-0120 OSLO
Norway

☐ This person is also inventor.

Telephone No. +47 23 23 84 00

Facsimile No. +47 23 23 84 01

Teleprinter No.

State *(that is, country)* of nationality: NO

State *(that is, country)* of residence: NO

This person is applicant for the purposes of: ☐ all designated States ☒ all designated States except the United States of America ☐ the United States of America only ☐ the States indicated in the Supplemental Box

Box No. III FURTHER APPLICANT(S) AND/OR (FURTHER) INVENTOR(S)

Name and address: *(Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)*

Svingen, Børge
Borkdalen 6
N-7550 HOMMELVIK
Norway

This person is:

☐ applicant only

☒ applicant and inventor

☐ inventor only *(If this check-box is marked, do not fill in below.)*

State *(that is, country)* of nationality: NO

State *(that is, country)* of residence: NO

This person is applicant for the purposes of: ☐ all designated States ☐ all designated States except the United States of America ☒ the United States of America only ☐ the States indicated in the Supplemental Box

☒ Further applicants and/or (further) inventors are indicated on a continuation sheet.

Box No. IV AGENT OR COMMON REPRESENTATIVE; OR ADDRESS FOR CORRESPONDENCE

The person identified below is hereby/has been appointed to act on behalf of the applicant(s) before the competent International Authorities as: ☒ agent ☐ common representative

Name and address: *(Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country.)*

Leistad, Geirr I.
of Thin Film Electronics ASA
P.O.Box 1872 Vika
N-0124 OSLO
Norway

Telephone No. +47 23 23 84 40

Facsimile No. +47 23 23 84 41

Teleprinter No.

☐ Address for correspondence: Mark this check-box where no agent or common representative is/has been appointed and the space above is used instead to indicate a special address to which correspondence should be sent.

Continuation of Box No. III FURTHER APPLICANT(S) AND/OR (FURTHER) INVENTOR(S)

If none of the following sub-boxes is used, this sheet should not be included in the request.

Name and address: *(Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)*

Risvik, Knut Magne
Nygata 11
N-7014 TRONDHEIM
Norway

This person is:

- [ ] applicant only
- [x] applicant and inventor
- [ ] inventor only *(If this check-box is marked, do not fill in below.)*

State *(that is, country)* of nationality: NO

State *(that is, country)* of residence: NO

This person is applicant for the purposes of:

- [ ] all designated States
- [ ] all designated States except the United States of America
- [x] the United States of America only
- [ ] the States indicated in the Supplemental Box

Name and address: *(Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)*

Halaas, Arne
Overlege Bratts vei 60
N-7026 TRONDHEIM
Norway

This person is:

- [ ] applicant only
- [x] applicant and inventor
- [ ] inventor only *(If this check-box is marked, do not fill in below.)*

State *(that is, country)* of nationality: NO

State *(that is, country)* of residence: NO

This person is applicant for the purposes of:

- [ ] all designated States
- [ ] all designated States except the United States of America
- [x] the United States of America only
- [ ] the States indicated in the Supplemental Box

Name and address: *(Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)*

Egge, Tor
Hjalmar Haalkes vei 1
N-7033 TRONDHEIM
Norway

This person is:

- [ ] applicant only
- [x] applicant and inventor
- [ ] inventor only *(If this check-box is marked, do not fill in below.)*

State *(that is, country)* of nationality: NO

State *(that is, country)* of residence: NO

This person is applicant for the purposes of:

- [ ] all designated States
- [ ] all designated States except the United States of America
- [x] the United States of America only
- [ ] the States indicated in the Supplemental Box

Name and address: *(Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)*

This person is:

- [ ] applicant only
- [ ] applicant and inventor
- [ ] inventor only *(If this check-box is marked, do not fill in below.)*

State *(that is, country)* of nationality:

State *(that is, country)* of residence:

This person is applicant for the purposes of:

- [ ] all designated States
- [ ] all designated States except the United States of America
- [ ] the United States of America only
- [ ] the States indicated in the Supplemental Box

- [ ] Further applicants and/or (further) inventors are indicated on another continuation sheet.

Box No.V DESIGNATION OF STATES

The following designations are hereby made under Rule 4.9(a) *(mark the applicable check-boxes: at least one must be marked)*:

Regional Patent

☒ AP **ARIPO Patent:** GH Ghana, GM Gambia, KE Kenya, LS Lesotho, MW Malawi, SD Sudan, SL Sierra Leone, SZ Swaziland, TZ United Republic of Tanzania, UG Uganda, ZW Zimbabwe, and any other State which is a Contracting State of the Harare Protocol and of the PCT

☒ EA **Eurasian Patent:** AM Armenia, AZ Azerbaijan, BY Belarus, KG Kyrgyzstan, KZ Kazakhstan, MD Republic of Moldova, RU Russian Federation, TJ Tajikistan, TM Turkmenistan, and any other State which is a Contracting State of the Eurasian Patent Convention and of the PCT

☒ EP **European Patent:** AT Austria, BE Belgium, CH and LI Switzerland and Liechtenstein, CY Cyprus, DE Germany, DK Denmark, ES Spain, FI Finland, FR France, GB United Kingdom, GR Greece, IE Ireland, IT Italy, LU Luxembourg, MC Monaco, NL Netherlands, PT Portugal, SE Sweden, and any other State which is a Contracting State of the European Patent Convention and of the PCT

☒ OA **OAPI Patent:** BF Burkina Faso, BJ Benin, CF Central African Republic, CG Congo, CI Côte d'Ivoire, CM Cameroon, GA Gabon, GN Guinea, GW Guinea-Bissau, ML Mali, MR Mauritania, NE Niger, SN Senegal, TD Chad, TG Togo, and any other State which is a member State of OAPI and a Contracting State of the PCT *(if other kind of protection or treatment desired, specify on dotted line)* . . . .

National Patent *(if other kind of protection or treatment desired, specify on dotted line)*:

☒ AE United Arab Emirates
☒ AL Albania . . . .
☒ AM Armenia . . . .
☒ AT Austria . . . .
☒ AU Australia . . . .
☒ AZ Azerbaijan
☒ BA Bosnia and Herzegovina . . . .
☒ BB Barbados
☒ BG Bulgaria . . . .
☒ BR Brazil . . . .
☒ BY Belarus . . . .
☒ CA Canada
☒ CH and LI Switzerland and Liechtenstein
☒ CN China . . . .
☒ CR Costa Rica . . . .
☒ CU Cuba . . . .
☒ CZ Czech Republic . . . .
☒ DE Germany . . . .
☒ DK Denmark . . . .
☒ DM Dominica
☒ EE Estonia . . . .
☒ ES Spain . . . .
☒ FI Finland . . . .
☒ GB United Kingdom
☒ GD Grenada
☒ GE Georgia . . . .
☐ GH Ghana . . . .
☐ GM Gambia
☒ HR Croatia . . . .
☒ HU Hungary . . . .
☒ ID Indonesia
☒ IL Israel . . . .
☒ IN India . . . .
☒ IS Iceland
☒ JP Japan . . . .
☒ KE Kenya . . . .
☒ KG Kyrgyzstan . . . .
☒ KP Democratic People's Republic of Korea . . . .
☒ KR Republic of Korea . . . .
☒ KZ Kazakhstan . . . .
☒ LC Saint Lucia
☒ LK Sri Lanka
☒ LR Liberia
☒ LS Lesotho . . . .
☒ LT Lithuania
☒ LU Luxembourg
☒ LV Latvia
☒ MA Morocco . . . .
☒ MD Republic of Moldova . . . .
☒ MG Madagascar . . . .
☒ MK The former Yugoslav Republic of Macedonia . . . .
☒ MN Mongolia
☒ MW Malawi . . . .
☒ MX Mexico . . . .
☒ NO Norway
☒ NZ New Zealand . . . .
☒ PL Poland . . . .
☒ PT Portugal . . . .
☒ RO Romania
☒ RU Russian Federation . . . .
☒ SD Sudan
☒ SE Sweden
☒ SG Singapore
☒ SI Slovenia . . . .
☒ SK Slovakia . . . .
☒ SL Sierra Leone . . . .
☒ TJ Tajikistan . . . .
☒ TM Turkmenistan . . . .
☒ TR Turkey . . . .
☒ TT Trinidad and Tobago . . . .
☒ TZ United Republic of Tanzania
☒ UA Ukraine . . . .
☒ UG Uganda . . . .
☒ US United States of America . . . .
☒ UZ Uzbekistan . . . .
☒ VN Viet Nam . . . .
☒ YU Yugoslavia . . . .
☒ ZA South Africa . . . .
☒ ZW Zimbabwe . . . .

Check-boxes reserved for designating States which have become party to the PCT after issuance of this sheet:

☒ DZ Algeria . . . .
☐ . . . .

Precautionary Designation Statement: In addition to the designations made above, the applicant also makes under Rule 4.9(b) all other designations which would be permitted under the PCT except any designation(s) indicated in the Supplemental Box as being excluded from the scope of this statement. The applicant declares that those additional designations are subject to confirmation and that any designation which is not confirmed before the expiration of 15 months from the priority date is to be regarded as withdrawn by the applicant at the expiration of that time limit. *(Confirmation (including fees) must reach the receiving Office within the 15-month time limit.)*

Box No. VI PRIORITY CLAIM ☐ Further priority claims are indicated in the Supplemental Box.

Filing date of earlier application *(day/month/year)*	Number of earlier application	Where earlier application is: national application: country	regional application:* regional Office	international application: receiving Office
item (1) 10 May 1999 (10.05.99)	19992269	NO		
item (2)				
item (3)				

☒ The receiving Office is requested to prepare and transmit to the International Bureau a certified copy of the earlier application(s) *(only if the earlier application was filed with the Office which for the purposes of the present international application is the receiving Office)* identified above as item(s): (1)

* *Where the earlier application is an ARIPO application, it is mandatory to indicate in the Supplemental Box at least one country party to the Paris Convention for the Protection of Industrial Property for which that earlier application was filed (Rule 4.10(b)(ii)). See Supplemental Box.*

Box No. VII INTERNATIONAL SEARCHING AUTHORITY

Choice of International Searching Authority (ISA) *(if two or more International Searching Authorities are competent to carry out the international search, indicate the Authority chosen; the two-letter code may be used)*:

ISA / SE

Request to use results of earlier search; reference to that search *(if an earlier search has been carried out by or requested from the International Searching Authority)*:

Date *(day/month/year)* Number Country *(or regional Office)*

Box No. VIII CHECK LIST; LANGUAGE OF FILING

This international application contains the following **number of sheets**:

request	:	4
description (excluding sequence listing part)	:	11
claims	:	3
abstract	:	1
drawings	:	4
sequence listing part of description	:	
Total number of sheets	:	23

This international application is **accompanied** by the item(s) marked below:

1. ☒ fee calculation sheet
2. ☐ separate signed power of attorney
3. ☐ copy of general power of attorney; reference number, if any:
4. ☐ statement explaining lack of signature
5. ☐ priority document(s) identified in Box No. VI as item(s):
6. ☐ translation of international application into (language):
7. ☐ separate indications concerning deposited microorganism or other biological material
8. ☐ nucleotide and/or amino acid sequence listing in computer readable form
9. ☐ other *(specify)*:

Figure of the drawings which should accompany the abstract: Figs 8, 9

Language of filing of the international application: English

Box No. IX SIGNATURE OF APPLICANT OR AGENT

Next to each signature, indicate the name of the person signing and the capacity in which the person signs (if such capacity is not obvious from reading the request).

Geirr I. Leistad

For receiving Office use only

1. Date of actual receipt of the purported international application:
2. Drawings: ☐ received: ☐ not received:
3. Corrected date of actual receipt due to later but timely received papers or drawings completing the purported international application:
4. Date of timely receipt of the required corrections under PCT Article 11(2):
5. International Searching Authority (if two or more are competent): ISA /
6. ☐ Transmittal of search copy delayed until search fee is paid.

For International Bureau use only

Date of receipt of the record copy by the International Bureau: